Open Interconnect for the Internet Community (OpenNet)

Martin Potts[1], Edmundo Monteiro[2], Donal Morris[3]
Michel Diaz[4], Philip Jacobs[5], Jean-Marc Uze[6]

[1]Martel GmbH, Bern, Switzerland
martin.potts@martel-consulting.ch
[2]University of Coimbra, Coimbra, Portugal
edmundo@dei.uc.pt
[3]Red Zinc Limited, Dublin, Ireland
dmorris@redzinc.net
[4]LAAS/CNRS, Toulouse, France
Michel.Diaz@laas.fr
[5]Cisco Systems International
phjacobs@cisco.com
[6]Juniper Systems EURL
juze@juniper.net

Abstract. This paper describes the OpenNet FP6 IST Specific Support Action (SSA) project. OpenNet addresses one of the main barriers to the wider usage of the Internet for all services; namely, the lack of predictable support for aspects such as end-to-end quality of service in a broader sense (availability, throughput, delay, delay variation, loss, security and privacy) when packets have to cross domains operated by different providers. As the need for a new business model for the Internet grows, the pace of investigation into this topic is increasing. The objective of OpenNet is to engage Europe more strongly in the discussions towards finding a solution. A particular development that will be followed actively is the definition of a business layer, as is being worked on in the worldwide IPSphere Forum. This layer is intended to manage all the business aspects (policies, pricing, SLA, service discovery, etc.) in order to enable on-demand secured and assured experience to the application, from end-to-end across multiple operators/domains.

Keywords: End-to-end QoS, Reliability, Security, IPSphere

1. Introduction

The OpenNet (Open Interconnect for the Internet Community) IST FP6 (Information Science Technologies, 6th Framework Programme) European Project addresses one of the main barriers to the wider usage of the Internet for all services; namely, the lack of predictable support for end-to-end quality of service (QoS) in a broad sense including aspects like: availability, throughput, delay, delay variation, loss, security and privacy, when packets have to cross domains operated by different providers.

Whilst various - sometimes proprietary - solutions are available within domains, the only current industry-accepted approach used to enable the end-to-end delivery of packets with characteristics that are needed for services such as Triple Play and VPNs, is the static and inflexible mechanism of interconnection contracts, known as

Service Level Agreements (SLAs). The parameters within these SLAs are also generally not adequate to specify the requirements for a specific service, providing only the characteristics at a high level of aggregation (e.g. for a link as a whole).

As the need for a new business model for the Internet grows, the pace of investigation into this topic is increasing, and the objective of OpenNet is to involve Europe more strongly in the discussions towards finding a solution. A particular development that will be followed actively is the definition of a business layer, as (for example) is being worked on in the worldwide IPSphere Forum. This layer is intended to manage all the business aspects (policies, pricing, SLA, service discovery, etc.) in order to enable on-demand secured and assured experience to the application, from end-to-end across multiple operators/domains. Underneath this new business layer, providers will be free to exercise whatever IP business model they choose, including the current Internet model, but they could also adopt whatever network-layer standards made sense to them and to their equipment vendors.

The technical approach is to bring together the major Internet router manufacturers worldwide (Cisco, Juniper), GÉANT (which has activities on BoD, End-to-end, SLA Monitoring, etc.), NRENs, IST projects and the IPSphere Forum, through the organisation of workshops. The outcome will be consensus towards the definition of this business layer, including the agreement of the parameters and values that are appropriate for the multiple services that are carried across the inter-domain interfaces.

The benefit will be the avoidance of the current situation where every ISP has to make SLAs with its peers, which will often be using incompatible bundles of QoS definitions for delay, delay variation and loss. This existing scheme is time-consuming, inefficient in terms of both capital and operational expenditure, not scalable, and impossible to manage end to end, when several domains are crossed.

The output will be reflected in reports and Internet Drafts, and will complement recommendations on quality of service and class of service parameters and values from the ITU, such as G.1010 and Y.1541.

This paper is structured as follows: Section 2 presents the IPSphere model; Section 3 describes the OpenNet objectives and expected contributions; Section 4 presents VPNSphere, a project that is aiming to validate IPSphere concepts; Finally, Section 5 discusses the main conclusions.

2. The IPSphere Model

The telecommunications industry has embraced the Internet Protocol (IP) as the ubiquitous vehicle for the new worldwide communication infrastructures. This is evident in the movement of the new wireless and wire-line sectors towards a replacement of the PSTN, the growing support of new Internet services and also in the 3G and 4G wireless business.

As IP is moving towards the support technology for the future communications infrastructures, new requirements arise that impose new design principles to the Internet architecture. A central requirement is to allow market forces to express themselves for the end-user benefit [1], in a market model consisting of multiple

providers in a service delivery chain. To support this requirement the definition of standard interfaces between the market players and the inclusion of quality of service, reliability, trust and security ingredients in addition to the current IP service model is needed [2, 3, 4, 5].

The IPSphere Forum initiative proposes a new architecture for the development and integration of business requirements in today's Internet aiming at the development of an end-to-end service oriented Internet powered with QoS, reliability and security mechanisms. Some of the authors of the present paper are actively involved in the developments of this new business oriented architecture for the Internet. Figure 1 describes the reference architecture and the interfaces and stratums identified and currently being specified by the IPSphere Forum [6].

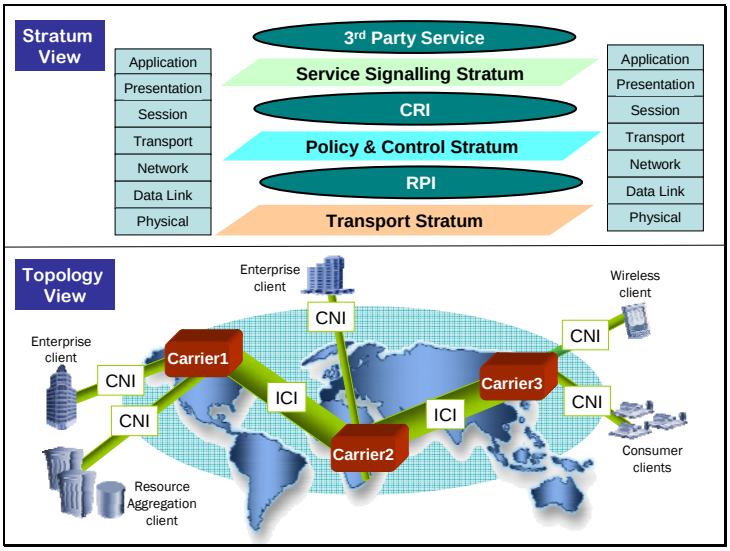

Figure 1: IPSphere Interfaces and Stratums [6]

The interfaces are: Client-Network Interface (CNI), Inter-Carrier Interface (ICI), Client Resource Interface (CRI) and Resource Provisioning Interface (RPI). The stratums are: service Signalling Stratum, Policy and Control Stratum and Transport Stratum. These interfaces and stratums will be described in detail in the following two subsections.

2.1 Client-Network Interface

The Client-Network Interface (CNI) enables applications to automatically request from the network the capabilities they require. A video-based application, for instance, could dynamically reserve resources that prevent disruptive jitter or delay. So the CNI consists of the signalling, data transfer between client and network in order to ensure authorized usage. It represents application's network requirements, and handles ingress/egress user's data traffic. Here the client means an application

'end-point' – 'originator' or 'terminator'. The CNI will handle four dimensions which are the bandwidth, quality, security and availability. Its flexibility allows to adapt to all scenarios such as service centres (richest interface, largest scale), wireless (undeterred, roaming capable), consumer (smaller scale, less but growing symmetry, less sophistication) and enterprise (smaller scale, symmetric traffic flow, privacy required).

In any scenario, CNI will receive application requests and adjust the service to meet specific needs. This new mechanism is required to allow users and applications to request from the network specific changes in quality, security, and availability. The network must be capable of understanding and then enforcing these requests. The CNI complements the federation concept outlined in the item above by automating and standardizing the way users and applications signal these requests to the network. As discussed previously, these requests allow the network to establish identity and membership for a particular service federation. Members of this federation may then signal the network for service additions, moves, and changes. Policing these requests also requires that IPSphere implementations provide stateful techniques for monitoring traffic flows and maintaining user identity. Stateful flow monitoring – both in-line and off-line – using standards such as *cflow* must be built into the network. Stateful privacy and security mechanisms such as NAT and stateful firewall are also key components of the CNI interface.

2.2 Inter-Carrier Interface

An Inter-Carrier Interface (ICI) enables one service provider's network to identify to another the level of service being provided and the network resources required to meet that service level. In addition, the ICI defines how the receiving service provider is compensated for handling that traffic.

Multi-national corporations, global content providers and regulatory laws are all reasons why the "ends" of end-to-end communication may not lie within a single carrier. In this case, the involved carriers must cooperate to deliver the required quality and to bill each other correctly. Also, every carrier must have a mean to see how well it (and its partners) are delivering the service. So the Inter-Carrier Interface (ICI) consists of the signalling between/across operator jurisdictions. It ensures performance of transiting traffic, to incorporate settlement between providers.

IPSphere implementations will enable services and applications built on the strengths of each domain to extend beyond the geographic boundaries of any one network operator or business entity. Meeting these criteria demands improved the standard – specifically, the development of an Inter-Carrier Interface (ICI) which allows network operators to exchange service metrics and settlement information.

To enable this final stage of evolution, the following two requirements must be met.

First, the support of QoS across multiple domains for end-to-end performance – the ultimate goal for the global project is an environment in which services and applications can be extended to anyone, anywhere – without the need for additional facilities. To accomplish this, standards will be required to ensure that all pertinent service information – such as QoS markings, security policy, and performance metrics

– are reconciled and maintained transparently when a user's traffic crosses service provider boundaries. However, the ICI not only sets the grounds for more compelling service offerings, it also paves the way for greater operational convergence for network operators. Service and management federations, as discussed above, provide an automated mechanism for identity and trust validation. In the global IPSphere powered Internet, enabled federations are also extended across geographical boundaries – providing trust validation and certification for inter-domain relations.

Also, it is necessary to communicate the state of one domain to other domains for congestion or failover scenarios – finally, the project must raise the bar considerably for availability and reliability. As the common infrastructure for all services and applications, IPSphere delivers mechanisms for communicating state at both the network or link level and at the service or user level. Only by binding these domains, can application and service availability be assured. At the network level, devices must provide mechanisms such as MPLS across network-to-network interfaces, RSVP, BFD, and GMPLS. At the service or user level, the network must provide not only the automated membership we have already outlined, but also automated responses to events such as attacks or network outages and general policy changes. In addition, network operators must be supplied with tools to measure availability, such as SLA measurement & verification.

Standard business practices and mechanisms – to ensure that services can truly be extended beyond a single provider or geography, the global IPSphere powered Internet must also provide tools and practices for exchanging and enforcing the business aspects of a service. For example, in-depth monitoring, troubleshooting, and reconciling are essential for accounting and proper settlement. While these are business- related tools, the technical requirements that underpin them are non-trivial. QoS, security, and availability parameters must all be exchanged and enforced, and monitoring these policies necessitates uncompromising performance at many levels.

2.3 IPSphere Stratums

Besides the interfaces described in the previous section the IPSphere model also defines three stratums to reflect the different levels of abstraction in the proposed architecture.

First, the Client-Resource Interface (CRI) specifies how the Service Signalling stratum and Policy and Control stratum interact. CRI is a new concept, and could be based on SOAP/XML.

Second, the Resource Provisioning Interface (RPI) specifies how the Policy and Control stratum and Transport stratum interact. In many current networks RPI-like protocols can be vendor-specific and proprietary, such as CLI. In addition, the RPI includes SNMP, XML, *syslog*, and accounting data, along with service provisioning protocols such as COPS-PR. It is an intention of the project to recommend a vendor-independent protocol for the RPI.

The two protocols described above make the link between the trust barrier state and the trusted state (or service operation state): initially, only service signalling is possible with the CNI, and no client-client data exchange is possible. Then the client makes claims, composed of the identity, the desired service(s) etc. If the claims are

validated, policy and control is signalled via CNI. Policy and Control stratum manage then the Packet Handling behaviour in associated network elements via RPI. The client-client data exchange is now possible, with packets treated as outlined in requested service description(s). The Service Signalling remains available, as needed for service modification or termination. CNI and RPI play a determinant role to move the communication from a trust barrier state (Service Signalling only) to a trusted state (Client to Client and Service Signalling).

3. The OpenNet Project

Today the Internet functions in two dimensions: connectivity and bandwidth (see Figure 2 below). Quality, reliability and security are missing. The results from OpenNet will provide the capability to offer assurances in multiple dimensions: quality, reliability and security, without compromising the connectivity and bandwidth dimensions.

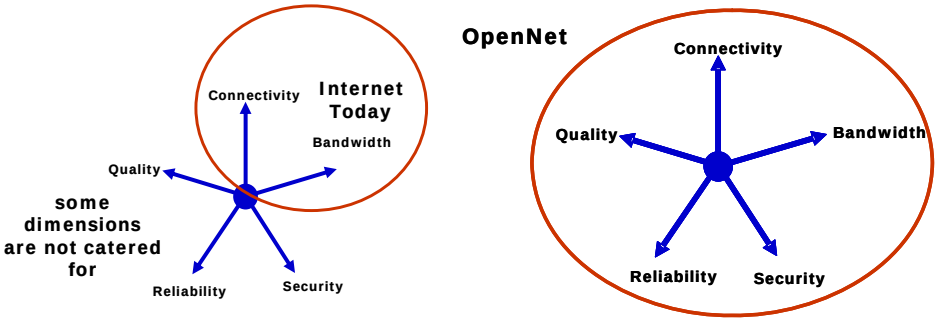

Figure 2: New dimension of OpenNet

The present Internet communication services do not offer more than a best-effort quality of service which is only sufficient in the context of an over-provisioned network. The current ISP model does not allow the ISP to guarantee the quality of real-time services because, for example, it has no view of the access topology and it is not notified when congestion occurs at that level. To take into account the constraints of real-time services (e.g. audio, video) the access and core network operators must have a role to play in the QoS resource allocation for the benefit of both end-users and service providers.

The new dimensions to be addressed by the OpenNet project will contribute to the final merging of the traditional telecommunication business with the Internet, and the deployment of a full range of new communication services. The following three problems represent the root of the Internet profit dilemma:

- The fixed-price model of Internet access makes it difficult to charge more for new uses of the Internet. This has encouraged the exploitation and expansion of Internet usage, but without a proportional increase in the revenues of those who must invest to expand the capacity of the Internet;

- The lack of commercial-grade interfaces between ISPs - i.e. interfaces that include guarantees of performance and settlement for service participation - has discouraged the deployment of enhanced features like QoS. Technical means exist to provide QoS between ISPs, but there is no legal requirement for inter-ISP settlement (like that for most PSTN common carriers), and no business-layer solution from the Internet community;
- Adding capacity to an IP network means adding resources. While those resources can be offered for sale as "priority services" supporting a theoretical premium service class, they are generally used just to make "best effort" services better, and bring no additional revenue.

In addition, the requirements of user communities are becoming evermore demanding of QoS, to the extent that bandwidth alone will not solve the problem. The OpenNet project will cooperate in the IPSphere Forum in the development of the interfaces and corresponding architectural framework that will have the appropriate data models and support mechanisms in terms of trust and security, quality of service and reliability, so that carriers and service providers can deploy them in the emerging scenario of multi-provider service independent heterogeneous network infrastructure.

The quest for the new generation Internet architecture attracted a number of initiatives (see Figure 3 below), in addition to the traditional standards bodies and the numerous forums, to foster the development of the Internet as it emerges into this new world infrastructure.

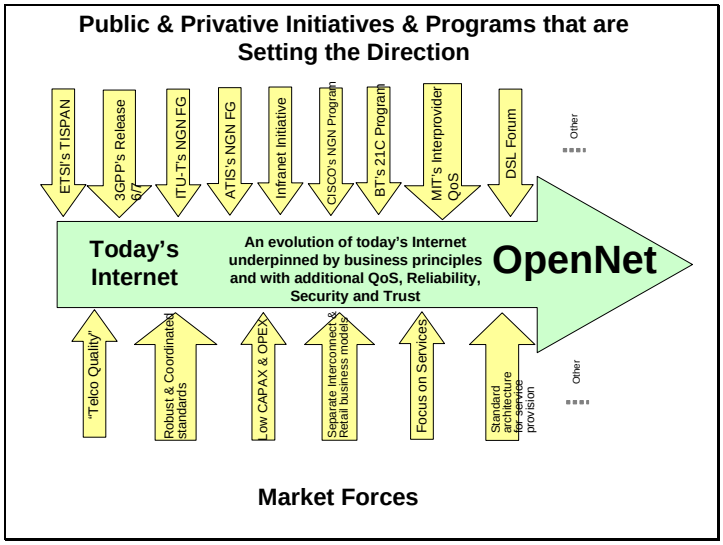

Figure 3: Initiatives and Programs driving OpenNet

The new generation internet initiatives consisting of both public and private activities, including New-Arch [7], 21st Century networking (21C) [8], IP-NGN Program [9] and the IPSphere Initiative [6], are seeking to foster models, working

methodologies, standards and components for the new era infrastructure. They all and especially IPSphere will be taken into consideration for OpenNet.

We can think of networking as having two dimensions: the application experience (e.g. a user accessing a content server to watch a movie) and the network transport (e.g. the IP networks built all over the world). Today, these two dimensions are very much separate – this is a problem. The risk in this scenario is that the network is reduced to "plumbing" and all the value and intelligence resides in the application dimension.

To combat this, OpenNet will promote new protocols – of which IPSphere CNI and ICI are examples – which bind or "couple" the application experience and the network transport. In this scenario, a user may now access the same content server to watch a movie, but be assured a particular level of bandwidth, quality and security. This is a fundamental difference. The application must now be aware of the network and vice-versa. This is a key distinction as this new model brings intelligence into the network and ensures that the network is an essential part of the overall application experience.

So how exactly does this coupling of the two worlds happen?

The simplest way to describe this process is to say that users, applications, and network resources become members of a closed community – this community is called a "federation". Once the appropriate entities are part of this federation, they have exactly the quality, bandwidth, security, and reliability necessary for their application experience – nothing more than this, and nothing less than this. This is a big difference from the public Internet where it is nearly impossible not to get more than you want (e.g. SPAM).

The key to build this federation is the establishment of identity and trust. A user or application requesting the use of a network application or resource must first cross a "trust barrier" through authentication and validation. Once that entity is identitied and the trust level is verified and set, the user or application is allowed into the federation and assured the appropriate level of service.

3.1 OpenNet Objectives

The OpenNet project has two main objectives.

First, to address the issue of enabling carriers and IP service providers to profit from the support of premium services, which will encourage them to provide more sophisticated and entertainment-oriented services over the Internet.

Second, to involve Europe more strongly in the discussions towards finding a solution for the definition of a business layer, as (for instance) is being worked on in the worldwide IPSphere Forum. This layer is intended to manage all the business aspects (e.g. policies, pricing, SLA, service discovery) in order to enable on-demand secured and assured experience to the application, from end-to-end across multiple operators/domains.

The benefit will be the avoidance of the current situation where every ISP has to make individual interconnection contracts (Service Level Agreements - SLAs) with its peers, which will often be using incompatible bundles of QoS definitions for delay, delay variation and loss. This existing scheme is time-consuming, inefficient in terms

of both capital and operational expenditure, not scalable, and impossible to use to manage end to end, when several domains are crossed.

3.2 OpenNet Technical Approach

As an SSA (Specific Support Action Project), OpenNet focuses on exchanging information and generating consensus. The key issues that will be studied are the interfaces that allow clients, access and core network operators to interoperate to obtain - and provide to their customers - a level of QoS. These interfaces will create a system context including the signalling, policy, security, reliability and QoS features to enable service differentiation for users, and competition between providers, in a worldwide, open telecommunications market. This market evolution will support a socio-economic evolution towards the availability of low cost and generalised broadband access through the delivery of new services.

The technical approach followed by the project is to bring together the major Internet router manufacturers worldwide (Cisco, Juniper), GÉANT (which has activities on BoD, End-to-end, SLA Monitoring, etc.), NRENs, IST projects and the IPSphere Forum, through the organisation of workshops.

OpenNet will take an integrated and comprehensive approach, by involving all relevant stakeholders of the QoS value chain, and addressing these issues at different levels and from different perspectives.

These workshops will ensure that there is a full exchange of knowledge and experiences between experts in the field, in order that all parties are aware of the requirements, capabilities and limitations of the various techniques being proposed by the Internet community.

The measurable outputs of the project will be the following:

- A strengthening of Europe's presence in this international collaboration activity;
- Consensus towards the definition of the business layer, including the agreement of the parameters that are appropriate for the multiple services that are carried across the inter-domain interfaces;
- Many FP6 projects have an interest in the support of QoS over IP networks (eg. MUSE, EuQoS, GN2, Eurolabs, EGEE). Opening this opportunity for partners to discuss the topic with the leading companies in the field will give them the chance to orientate their developments towards implementations that will be compatible with emerging equipment;
- Support for FP7: it is anticipated that areas for related future joint RTD activities in FP7 will be identified. These will be documented and discussed with the Commission.

The validation work that will be conducted by OpenNet is essential for preparing future upgrades of the Internet infrastructure deployed across Europe and for fostering the early deployment of Next Generation Information and Communications Networks. The work of this SSA is clearly complementary to - and in support of - the activities carried out in the area of Research Infrastructures on high-capacity and high-speed communication networks for all researchers in Europe (GÉANT) and to

high performance Grids, which represent major components of the *e-Infrastructure* concept. For example, OpenNet supports these testing and validation activities, by being the mechanism through which outputs from those S&T projects that have activities in the area of QoS (e.g. EuQoS, GN2, EGEE, Eurolabs) can be disseminated, whilst providing a channel for the transfer of information back into the Programme.

Through its series of workshops, backed up by the practical experience of the partners, OpenNet will "*… promote interoperability across heterogeneous technology domains, facilitate interoperability of solutions across different scientific and industrial communities, support the creation of standards, promote economies of scale during the validation phase and achieve broader-scale up-take of technology across numerous user communities.*"

4. Validation with VPNSphere

Besides the OpenNet evangelization initiative described above the project partners are also involved in an experimental project, named VPNsphere [11] to develop and validate some of the IPSphere concepts in a prototype environment

The IPSphere / OpenNet achievements in terms of architecture and interface definition will be validated in the VPNSphere scenario that include the most challenging services and environments for the interfaces to be developed and in terms of provisioning QoS, security and reliability. This validation scenario will be focused in the IP VPN (Virtual Private Network) use case.

From a business application perspective the contemporary Internet is simply not enabled to facilitate most of the highly valued corporate applications. These are still transported using private technologies such as Frame Relay, Private Leased Line, and ATM. The IP VPN market in Europe was 2.730 M€ in 2003 and estimate to 8.560 M€ in 2008 with 70% of its potential market [10]. These networks are by their nature constrained in reach and flexibility. The organizations and enterprises relying and making use of them are therefore unavoidably constrained in a similar way.

Furthermore, network service providers are forecasted to face a critical period over the next several years. In the near future, no longer the ability to provide connectivity, e.g. for voice calls, will be considered as essential network service provider capability, but the ability to sell value-added services such as QoS, Reliability, or Security on top of plain connectivity and data transportation. This is foreseen to be caused by driving factors like consumer data and content. Affected from this fundamental change are not only the network service provider revenues, but also their survival.

Layer 3 Multi-provider VPN use case will build a VPN platform model spanning a number of continents and providers. It will offer, in addition to a base connectivity, data service, a number of dynamic QoS based conversational voice and video services and security. The use case is mostly devoted to business customer offering office remote access for the employee with security, QoS and reliability.

The VPNSphere project will take three steps in formulating, developing and verifying the interfaces to be defined:

- First step: Research and specification of the architecture and framework, its interfaces and the metrics to be used to express QoS, reliability and security;
- Second step: Build the prototypes over Cisco and Juniper state of the art equipment to support the VPN use cases discussed above;
- Third step: Test and validate the functionality developed against the metrics defined in the first step.

The innovation of the OpenNet approach will be granted by the selection of the appropriate state-of-art ingredients to be integrated in the architecture and enhanced with the appropriate mechanisms to guarantee the required level of QoS, reliability, security and trust at the various interfaces. Some of the research topics needed for the VPNSphere project are currently being address by project partners in ongoing research projects.

The VPNSphere project started the 1st of January the partners are involved in the requirements phase. Results will be provided as soon as they are available.

5. Conclusions

This paper describes the OpenNet IST FP6 SSA project that was recently approved by the European Commission.

The OpenNet project aims to cooperate with the IPSphere Forum architectural and standardization initiative at European level, helping to bring a business layer to the Next Generation Internet.

The OpenNet and IPSphere achievements in terms of architecture and interface definition are being validated in the VPNSphere project that is focused in the IP VPN use case scenario.

Preliminary results show that the IPSphere concepts can bring a new vitality to the already old best-effort Internet, enabling the wider deployment of profitable new end-to-end services to users with QoS, security and reliability requirements.

Acknowledgments

This work was partially funded by the European Union 6th Framework Programme under contract IST FP6 SSA 035185OpenNet Project.

References

[1] David Clark et al, Tussle in Cyberspace: Defining Tomorrow's Internet, ACM SIGCOMM 2002, Pittsburgh, 2002

[2] David Hutchison, "Risk and Resilience in Emerging Networks", Keynote presentation at QofIS 2004: Barcelona, October, 2004.

[3] Ian Akyildiz, “QoS Past, Present and Future”, Keynote presentation at QofIS 2004: Barcelona, October, 2004.
[4] Vint Cerf, “The Internet is for Everyone”, Internet Society, RFC 3271, 2002.
[5] United Nations, “Creation of a global culture of cyber security and the protection of critical information infrastructures”, Resolution n. 58/199, UN, 2003.
[6] IPSphere Forum,“IPSphere architecture”, available at www.ipsphereforum.org
[7] New-Arch Project, “Future-Generation Internet Architecture” Available at www.isi.edu/newarch
[8] British Telecom, “21CN Reference”, available at www.bt.co.uk
[9] Cisco, “Cisco IP-NGN Program”, available at www.lightreading.com
[10] Frost & Sullivan, “P VPNs”, available at www.frost.com
[11] VPNsphere, “VPNsphere Project”, available at www.ist-newnet.org